



Applications *Run Faster &  
Perform More Consistently*  
on our Storage!

**Costa Hasapopoulos**

Chief Field Technology Officer &  
WW Vice President of Business Development

[www.pavilion.io](http://www.pavilion.io)

# WHY We Exist - Applications Run Faster on Pavilion...

## Databases (SQL or NoSQL)

---



## External Scale-Out File Systems

---



## Virtualized Infrastructure

---



## AI & Machine Learning

---




GPUDirect Storage  
Block and File

# What We Deliver – Predictable and Consistent Performance and Scale

Single Storage Platform Unit = 4 RU up to 2.2PB

Flexibility to support **multiple STORAGE types** on a single platform



20M IOPS Read

5M IOPS Write

**100 µs Read Latency**

**25 µs Write Latency**

120 GB/s Read

90 GB/s Write



**Shared Global Namespace**

78 GB/s Read

56 GB/s Write

52 GB/s Read

28 GB/s Write

File and Object **Shared Global Namespace** can be clustered with multiple systems to linearly scale performance and capacity

Not just high **performance READ** but also **high performance WRITE!**

# WOW the Results



**2x+** the Queries per Second (QPS) and  
**7x+** the Transactions per Second (TPS) at  
**<50%** the Latency

**1.8x** the Queries per Second (QPS) and  
**6x** the Transactions per Second (TPS) at  
**<60%** the Latency



**1.5x** the Queries per Second (QPS) and  
**67%** the Latency

**1.13x** the Queries per Second (QPS) and  
**89%** the Latency



**2.3x** the Transaction per Minute (TPM) and  
**2.3x** the Operations per Minute (NOPM)

**2x** the Transaction per Minute (TPM) and  
**1.3x** the Operations per Minute (NOPM)



elasticsearch

**1.10x** the Documents / Second  
**1.47x** the Operations / Second

**1x** the Documents / Second  
**.47x** Operations / Second



**1.14x** the Operations / Second  
**77%** the Read Latency  
**76%** the Update Latency

**1.06x** the Operations / Second  
**94%** the Read Latency  
**101%** the Update Latency



**< 200**  $\mu$ sec Latency for  
Reads & Write

**Similar or Same**  
Performance as Bare Metal



Not STAC benchmarks

Pavilion Data Systems. ©2022 All rights reserved.

# How do we do it?


Thank you!

Not **2** controllers – **20** controllers

Not **4** 32Gb FC or 50GbE – **40** x 100 GbE/EDR

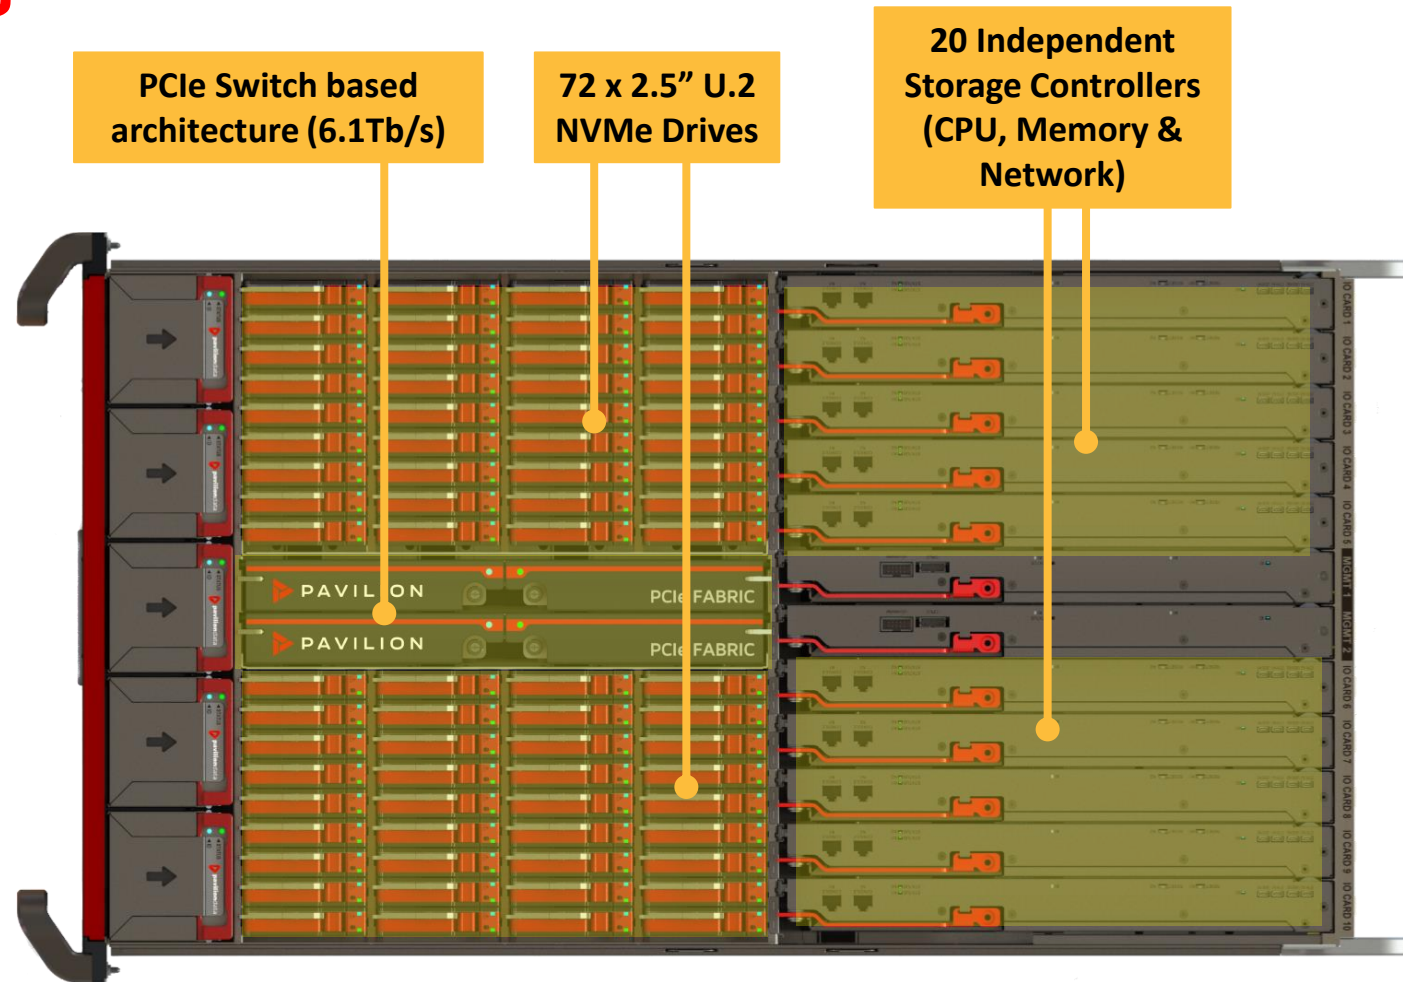
Not just **BLOCK** but also **FILE** & **OBJECT**

Not **RACKS** of equipment but **4 Rack Units**



4RU Chassis

40 x 100 GbE/EDR or  
10 x 200 HDR  
InfiniBand ports



PCIe Switch based  
architecture (6.1Tb/s)

72 x 2.5" U.2  
NVMe Drives

20 Independent  
Storage Controllers  
(CPU, Memory &  
Network)

Unrivalled Performance Density



Thank You!